



GB99/01772

The
Patent
Office

PGT/GB99 3 / 0 1 / 1 4



REC'D 30 JUL 1999
WIPO PGT

INVESTOR IN PEOPLE

**PRIORITY
DOCUMENT**
SUBMITTED OR TRANSMITTED IN
COMPLIANCE WITH RULE 17.1(a) OR (b)


The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.



Signed

Dated

28 JUN 1999

THIS PAGE BLANK (USPTO)

Patents Act 1977
(Rule 16)

29 JAN 1999

RECEIVED 29 JAN 1999
TELETYPE UNIT - PATENT OFFICE

The Patent Office

Cardiff Road

Newport

Gwent NP9 1RH

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

9902052.1

1. Your reference

A25745

2. Patent application number

(The Patent Office will fill in this part)

3. Full name, address and postcode of the or of each applicant (underline all surnames)

BRITISH TELECOMMUNICATIONS public limited company
81 NEWGATE STREET
LONDON, EC1A 7AJ, England
Registered in England: 1800000

Patents ADP number (*if you know it*)

1867002

If the applicant is a corporate body, give the country/state of its incorporation

UNITED KINGDOM

4. Title of the invention

COMMUNICATIONS NETWORK

5. Name of your agent (*if you have one*)

LIDBETTER, Timothy Guy Edwin

"Address for Service" in the United Kingdom to which all correspondence should be sent (*including the postcode*)

BT GROUP LEGAL SERVICES
INTELLECTUAL PROPERTY DEPARTMENT
HOLBORN CENTRE
120 HOLBORN
LONDON, EC1N 2TE

Patents ADP number (*if you know it*)

1867001

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (*if you know it*) the or each application number

Country

Priority application number
(*if you know it*)

Date of filing
(*day / month / year*)

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number of earlier application

Date of filing
(*day/month/year*)

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (*Answer 'Yes' if:*)

YES

- a) any applicant named in part 3 is not an inventor, or
- b) there is an inventor who is not named as an applicant, or
- c) any named applicant is a corporate body.

(See note (d))

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form

Description 27

Claim(s) 4

Abstract 1

Drawing(s) 15

10. If you are also filing any of the following, state how many against each item

Priority Documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (Patents Form 7/77)

Request for preliminary examination and search (Patents Form 9/77)

Request for substantive examination (Patents Form 10/77)

Any other documents (please specify)

11.

I/We request the grant of a patent on the basis of this application.

Signature(s)

Date:



29 January 1999

Timothy Guy Edwin LIDBETTER, Authorised Signatory

12. Name and daytime telephone number of person to contact in the United Kingdom

Lisa Colton

0171 492 8146

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 0645 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- Once you have filled in the form you must remember to sign and date it.
- For details of the fee and ways to pay please contact the Patent Office.

COMMUNICATIONS NETWORK

The present invention relates to a communications network, and in particular to charging mechanisms in such a network. It includes aspects of the inventions disclosed and claimed in the present applicant's co-pending British patent application no. 9812161.9 filed 5 June 1998 and the contents of that earlier application are incorporated herein by reference.

In conventional communications networks, such as national PSTNs (public switched telephone networks), a significant proportion of the network resources are devoted to metering and billing network usage. Studies have estimated these resources as consuming as much as 6% of the revenue of a telecommunications company. The Internet, by contrast, does not in general incorporate metering and billing mechanisms for individual customers. The absence of the network infrastructure required to support metering and billing reduces the operational costs of the Internet compared to conventional telephony networks, and has facilitated the rapid expansion of the Internet. However the absence of appropriate billing mechanisms has significant disadvantages in terms of the characteristics of the traffic carried by the internet. It encourages profligate use of network resources, and diminishes the incentive for investment in network infrastructure to support new applications requiring, e.g., guaranteed quality of service (QoS) and led to subscription based Internet access services.

According to a first aspect of the present invention, there is provided a method of operating a communications network comprising:

- a) measuring at each of a plurality of customer terminals usage by the the respective customer terminal of network resources; and
- b) subsequently calculating a network usage charge from the measurement data generated by step (a).

The present inventors have found that a key step in implementing a lightweight charging protocol suitable for use in a federated network is to decentralise the metering of network usage by arranging for each customer terminal to monitor its own use of network resources. In this way a charging mechanism is provided that is intrinsically scaleable and that avoids significant overheads within the network.

Preferably the method includes storing the measurement data generated by step (a). Preferably there is stored with the measurement data data identifying a tariff applicable to the said measurement data. The said data identifying the tariff may be the tariff itself, or may take the form of some identifying code or pointer for the tariff. Storing the tariff enables accounting data to be generated from measurements at the customer terminal even if the tariff varies over time.

Preferably the method includes communicating data generated by step (a) to a network accounting object controlled by a network operator. Alternatively data may be communicated from the network operator to the customer in a conventional way. The network usage data may be communicated explicitly and the charge for network usage calculated by the network operator. Alternatively the usage data may be communicated implicitly in accounting data indicating a charge calculated by the customer terminal.

Preferably the method includes a step carried out by the network operator of sampling part only of the traffic communicated between a customer terminal and the network. This sampled traffic is then compared with the network usage data reported from the customer terminal to the network provider accounting object, thereby detecting any discrepancy. The comparison may be of the total charged for network usage, or may be of the detailed measurement data. The former may be the norm for efficiency, with the latter used, in this case, only if the former shows discrepancies, in order to store evidence of fraud.

The inventors have found that the efficiency of the charging process can be further enhanced if the customer is responsible for measuring usage and providing usage data or priced usage data and the network operator measures only a sample of the customer traffic, on a random basis, to confirm the reliability of data provided by the customer.

Preferably the network operator accounting object is configurable to receive data either from a measurement object controlled by the network operator or from a customer terminal. Preferably the method includes changing from one configuration to the other in response to a control signal received at the network accounting object.

Preferably the method includes communicating measurement data to a system remote from the customer terminal. For example, data may be communicated from a number of customer terminals to a corporate accounting

system. The data may be sent explicitly, and/or a usage charge calculated using the data may be sent to the remote system. When data is reported to a remote system, this may be done immediately the data is generated, or may be done in the form of a report aggregating data from a series of measurements over a period
5 of time.

Preferably the method includes:

- communicating traffic between a customer terminal and a first network domain connected to the customer terminal,
- 10 further communicating the said traffic between the first network domain and a second network domain connected to the first network domain;
- communicating network usage data from the customer terminal to a first network accounting object in the first domain;
- communicating accounting data between the first network accounting object and a second network accounting object in the second domain.

15 This aspect provides a powerful and efficient method of accounting between domains in a federated data network. Although data may be flowing e.g from a first customer terminal, via intermediate network domains to a second customer terminal, the accounting data (i.e. the measurement data or data derived therefrom) need not all flow in the same direction. The invention encompasses, for
20 example, systems in which accounting data is passed from the customer to the first domain and also is passed from the second network domain to the first network domain.

Preferably the method includes determining from a current routing table in the first network domain the identity of a second domain communicating data with
25 the customer terminal via the first network domain, and communicating accounting data for the customer terminal with the second domain identified by the current routing table.

According to another aspect of the present invention, there is provided a method of operating a network comprising a plurality of network domains,
30 including calculating a charge for use by a respective customer of network resources, and making payment in settlement of the said charge to a third party clearer.

The invention also encompasses communications networks arranged to operate by the methods of the invention, and customer terminals, and network accounting servers for use in such a network.

Systems embodying the present invention will now be described in further detail, by way of example only, with reference to the accompanying drawings, in which:

Figure 1 is a schematic showing a network embodying the invention;

Figures 2a and 2b are schematics showing the component objects of a charging architecture for use with the network of Figure 1;

Figure 3a and 3b show data passed between the accounting objects of Figure 2a;

Figure 4 is a schematic showing protocol stacks on a customer terminal and in the network domain;

Figures 5a to 5e are class diagrams for software implementing accounting and measurement objects;

Figure 6 is a diagram showing a graphic user interface (GUI) for use with the objects of Figures 5a to 5e;

Figure 7 is a diagram showing the interface between neighbouring domains of the network of Figure 1;

Figure 8 is a diagram showing schematically the distribution of accounting data through multiple network domains;

Figure 9 is a diagram showing a network using service provider clearing;

Figure 10 is a diagram showing a network using third party clearing

As shown in Figure 1, a communications network 1 includes a number of network sub-domains 2A-C. The network sub-domains may be under the control of different operators. The operation of the network does not assume that there is mutual trust between the different operators. The network subdomains are interconnected by gateway routers 3, 4. In the present example the communications network is the Internet and supports both unicast and multicast Internet Protocol (IP) and associated protocols. A customer terminal 5 is connected via a public switched telephony network (PSTN) 6 and an access router 7 to a subdomain 2A. A single blocking test is applied to traffic at this point of access. The gateway routers 3,4, and access router 7 may be commercially available devices such as CISCO series 7500 routers and CISCO series AS5800

universal access server respectively. Other customer terminals are connected to the network, including a Java-enabled mobile terminal 8 and a data server 9.

The customer terminal 5 may be connected via a LAN to an accounting server. The accounting server may include an accounting object as described below that receives measurement data from the customer terminal.

Tariffs for the use of network resources are multicast through the network to the customer terminals. These tariffs are divided into bands of different volatilities. The tariffs are varied under the control of the network operators to reflect the overall loading of the network. That is to say, if network loading becomes high then the tariffs may be increased to reflect the scarcity of network resources. A network management platform 10 is connected to each subdomain. Each network management platform may comprise, for example, a computing system comprising a SPARC workstation running UNIX (Solaris) together with network management applications. The network management platform 10 hosts management entities and tariff entities. It may also function as an accounting server hosting network accounting objects as described below. The network management platform communicates with agents 100 in managed devices connected to the respective subdomain, for example using SNMP (simple network management protocol). The management platform monitors the overall loading of network resources in the respective subdomains, and adjusts the tariffs for network use accordingly. The Net management platform (NMP) instructs the agent to monitor the device and report aggregated results at regular intervals back to the NMP, so the NMP can monitor the combination of all reports.

In addition to this central control of the tariffs, a tariff algorithm at each customer terminal may be arranged to respond automatically to a locally detected variation in the loading of network resources. The use of local tariff variation is described and claimed in the present Applicant's co-pending application also entitled "Communications Network", BT reference A25626.

In the present example, charging is carried out using a "pay and display" process but traditional payment methods can alternatively be used. .. Figures 2a and 2b show the objects used to implement the charging architecture in this case. Figure 2a shows the higher level objects and 2b shows the component objects used in a software implementation of the architecture of Figure 2a and expands further the distribution of the accounting objects within a single domain. In Figure

2a, objects on the client terminal are shown in the half of the Figure labelled "customer" and objects on the access router 7 and the corresponding network sub-domain are shown in the half of the Figure labelled "edge network". The objects on the customer terminal include a session control object S, a customer business rules object B_c , a customer pricing object Pr_c , a QoS manager Q, a customer accounting object Act_c and a customer measurement object M_c . The business rules object B_c receives information on those aspects of the session which involve liability for payment and receives current pricing data from the pricing object Pr_c . The customer business object makes decisions, under the customer's policy control on which chargeable services are utilised, and how much of the chargeable services are utilised. These decisions are fed to the QoS manager Q, which decides which mechanisms are used to achieve the requirements. The QoS manager (and the accounting object) then controls the customer measurement object M_c to determine which aspects of traffic and service to measure and which aspects to ignore. The measurement object then records the selected aspects of the traffic, for example counting the number of packets transmitted and received by the customer terminal and the QoS levels for those packets. These data, together with the current tariffs, including any premium for congestion, are then used by the customer terminal to determine the charge payable to the network operator. The measurement object M_c is also programmed, by the accounting object, with instructions that determine the frequency at which data is reported to the customer accounting object Act_c . The customer accounting object Act_c passes accounting information (priced or not) to an accounting object Act_p in the network provider's domain. On the network provider's side, that is to say within the subdomain to which the customer terminal is connected, the customer's traffic is measured by a version of M, denoted M_p , but only on a sampling basis determined by the policing function, Po. That is to say, the network operator samples the customer's traffic only intermittently. Po controls where in the network measurements are made in order to capture all of any particular customer's traffic. A bulk measurement function, M_b , is responsible for reporting aggregate traffic levels, as reflected in the moving average of the router queue lengths, to the pricing object, Pr_p . Bulk measurements would typically be collected from across the provider's domain to a centralised pricing function (which would be replicated for reliability). Pr_p sets prices taking into account the

business rules from the network provider's business object, B_p , as well as the current traffic levels reported by M_b and pricing from neighbouring providers. The policing function, P_o , compares sample measurements from M_p with accounting messages received at Act_p as a result of the customers own measurements . If it

5 establishes that the accounts are insufficient it might restrict service at the access control gateway, Acs , or initiate some other punishment. Encapsulated within the accounting object another policing object checks the accounts match the payments within the contracted time for payment. Finally, the identity mapping function, I , provides a mapping between a customer's identity (account, digital

10 signature, etc.) and their current network address (typically allocated by the ISP, whether unicast or multicast).

The measurement (M) objects provide to the accounting (Act) objects the information that is required to create firstly accounting records and subsequently reports and bills. Measurement records are not stored as such in the Act objects:

15 measurement data is translated into accounting records as soon as possible. The translation of measurement data into accounting records involves a change of class type and some aggregation. In addition the measurement data may be linked to tariff information. The measurement data returned by the measurement objects includes, in this example, the following elements:

20 IP addresses of the two endpoints involved in the communication. This is readily available from the network packets.

Port numbers: These are used to distinguish between different services used by a user at one time. The port numbers are also available from the network packets.

Type of packets: service identity. This identifies the type of service, e.g. as RSVP,

25 as differential service or as data. This information allows different tariffs to be applied depending on the packet type.

Network usage information. This is the measurement data itself and may comprise, for example, a count of the number of packets.

Time period information. This, if element, when used, indicates the length of time

30 over which the measurement was made

Time reference. This may include a start time and an end time and may be used, for example, for applying discounts to traffic during defined "off-peak" hours.

In the presently preferred implementation, measurement data is returned by the measurement object to the Act object on an event-driven basis at time

intervals controlled by the accounting object. Alternative approaches may use polling of the measurement object by the Act object, or event driven polling. Communication of data may be effected using Java - RMI (remote method invocation) and the Java event model or a socket may be created between Act and M to send measurement objects. Further alternative communication mechanisms include the use of CORBA or SNMP like messaging. The present example makes use of an RMI/CORBA-like distributed event programming infrastructure called FLEXINET.

Measurement objects (M) offer a control interface to Act objects, so that Act objects can control what measures, and when and where M reports its measurement information. This control interface offers access to the following parameters:

1. Frequency at which measurement records are required (for a given customer or set of customers). This makes it possible to accommodate different accounting business models including, e.g., pay-as-you-go and traditional billing. The frequency may be specified as a period of a number of milliseconds.

2. What is to be reported to Act (for a given customer or set of customers). This parameter might specify all packets, or only packets with a give QoS threshold etc.

3. Where to report measurements (for a given customer or set of customers). This parameter may be a simple reference to the Act object or another business-related object for auditing or marketing purposes.

4. Current metering properties of the measurement object.

The Meter M at the network provider multiplexes the different measurement request for different customers and optimise the measurement and reporting processes.

The accounting objects on the customer terminal may be implemented using a small encrypted flat-file database. On the network provider's side, the equivalent objects may be implemented using a larger database that is scaleable to handle e.g., tens of thousands of customer accounts. An object request broker (ORB) is used for communication between the customer-side objects and the network-side objects, implemented using commercially available tools such as

ORBIX (TradeMark) from Iona Technologies plc. Serialisation is used to stream objects from one database to another via the network. The process of serialisation takes all the attributes of an object and streams the attributes over a specified medium together with information specifying the type of object that originated the data. A process of de-serialisation then takes the data from the transmission medium together with the object type information and creates a new object of the specified type and fills it with the data. The accounting databases hold a set of serialised accounting objects. The larger database required by the network provider may be an object-oriented database that accepts objects and serialises them into its storage space. Alternatively a non object oriented database may be used, in which case the accounting objects are translated into database types. For example the accounting objects are translated into SQL data types for use with a relational database.

The serialisation/de-serialisation mechanism described above is also used to support the measurement and accounting interface between network domains. For example, the edge-of-network domain that communicates packets to and from the customer terminal in turn passes packets to a number of neighbouring domains. Just as accounting data is passed from the customer to the edge-of-network domain, so also accounting data is passed from an accounting object 71 in the edge-of-network domain to an accounting object 72 in a neighbouring domain, and payment is made by the operator of the edge-of-network domain to the operator of the neighbouring domain. In this context, the edge-of-network domain is a retail domain, and the neighbouring domains are wholesale domains. As shown in Figure 7, the architecture of the interface between the retail domain and the wholesale domains is a recursive version of the interface between the retail domain and the end customer. However all the measurement and QoS features of the interface to the end customer are not required in the interface between the retail and wholesale networks. Where, as in this example, there are multiple wholesale providers, then the current routing and/or address allocation in the retail network is interrogated to apportion accounting between the wholesale networks. This is effectively another form of identity mapping, I. The mapping is needed between the identities of each neighbour provider and their current groups of unicast addresses, address prefixes, multicast addresses or autonomous system (AS) numbers. This is not generally required in the edge architecture, as an edge

customer typically has only one provider. If multiple providers were used by the customer, then mapping to apportion accounting is used at the edge too. As before, the measurement of traffic between retail and wholesale domains can be sampled and done in parallel to the data flow - no blocking is required. Any pair of
5 network providers might in practice each be mutual customers. In this case, the architecture for the retail/wholesale interface is mirrored so that all functions operate in both directions. Any payments between network domains are then determined by the balance of the products of each accounting flow and the relevant prices.

10 In a network comprising multiple domains then, as shown in Figure 8, a "wholesale" domain 82 may receive accounting data from a number of retail networks 81,83. These data are aggregated by the accounting object in domain 82 and then apportioned between further neighbouring domains, such as domain 84. The way in which the accounting data are apportioned is determined by an
15 averaged border routing table maintained in the domain 82. Figures 3a and 3b show the data which are passed between the accounting objects. In this example the account data comprises: account identity; bill record identity; service type identifier; source address; destination address; tariff identity; time; period (i.e. the period covered by the bill record); units; cost; and currency. In addition, the
20 payment data comprises the amount of money and the currency of payment.

Figure 4 shows the measurement region within protocol stacks on the customer terminal and in the retail network domain. Ideally there would be two measurement points within this region, one trusted by the customer and one trusted by the network, for example at the two points referenced (a) in the Figure.
25 For ease of implementation, a single measurement point (b) trusted by both parties may be used. This might be implemented, for example within a secure module such as a cryptographic card on the client terminal. As an alternative, measurements may be made at different points with some possibility of discrepancies between measurements. On the network the practical measurement
30 point is at the first access device(s) that, for each customer, inspects network layer headers (c)(IP in this case). ISPs should not measure any deeper into their network (d) because their access network and systems will introduce delays and losses.

For an individual customer (e.g. on dial-up access), a practical point at which to measure would also be alongside the network layer but in their end-system's stack (e). Ideally these measurement points would be lower in each stack to be closer to the interface between the two parties and less likely to be affected by contention in the stack. However, measuring at the link layer (f-f) would be inappropriate because only some chargeable parameters set at the network layer will ever be reflected in link layer frames; network level multicast, end-end latency requirements etc. may never be visible at the link layer. Also, link layer headers would need to be ignored when measuring packet sizes for bandwidth calculations to avoid apparent discrepancies where different link technologies are chained together.

In the reception direction (up the stack) this choice of measurement points implies that the lower layers must be dimensioned (buffer sizes, interrupt and thread scheduling priorities) to cope with the most stringent QoS requirements of higher layers. As frames are taken off the physical media, the machine must be able to pass data up the stack without any chance that usage-charged data gets discarded (e.g. due to buffer overflow caused by interrupt contention) before it gets to the network layer. It is at the network layer where the ISP's service is to be measured and where it is most convenient for QoS requirements to control correct differential treatment of the various flows as they are passed further up the stack (on end-systems) or forwarded (on routers).

The measurement objects described above may be implemented using, with appropriate modifications, publicly available network metering software such as Nevil Brownlee's NeTraMet system. This is a software meter which conforms to the IETF internet accounting architecture described in RFC 2063 and RFC 2064. The meter builds up, using "packet sniffing", packet and byte counts for traffic flows, which are defined by their end-point addresses. Although generally, Addresses can be ethernet addresses, protocol addresses (IP, DECnet, EtherTalk, IPX or CLNS) or 'transport' addresses (IP port numbers, etc), or any combination of these, in the present implementation IP addresses only are used. The traffic flows to be observed are specified by a set of rules, which are downloaded to NeTraMet by a 'manager' program. Traffic flow data is collected via SNMP (Simple Network Management Protocol) from a 'collector' program

Figures 5a to 5e are class diagrams illustrating an implementation of the measurement and accounting objects described above. The class diagrams are shown as a series of views.

The control view (5a) groups the classes related to control over the accounting class, including reporting control, metering-related control and general control functions. This view also relates to event dissemination. Control over the Accounting class is separated according to the type of control. This is why four interfaces are available. Two of those interfaces provide direct control over the behaviour of the Accounting object and the two others are related to a Java event model used to communicate both reporting information and measurement information. The ActControl interface provides control over the accounting class that relates to the accounting behaviour in general. It provides both methods to set a behaviour or properties and methods to find out about the current behaviour of the accounting object. For example, this interface is used to set the name of the accounting object or to query the Act object to find out a name previously given to the Act object. The ActReport interface provides control over issues related to account reporting. Control calls are directly related to the reporting behaviour of the accounting object. For example, a method named addReportListener() is used to register interest in reporting information. Once the registration is effective, subsequent calls to other control methods define behaviour such as the reporting frequency, request for immediate reporting, reporting security properties..etc. The two other listener interfaces (Report & Measurement) that the Accounting class implements are used to indicate that accounting objects are interested in accounting reports and measurements.

The accounting report view (Figure 5b) regroups the class related to the reporting behaviour and reporting process in the accounting objects. The accounting objects listens to accounting reports and generates such events as well. Accounting objects generate accounting reports and distribute them (using the traditional Java event model) to objects that have registered their interest in such events. In the present implementation flexinet (A CORBA like distributed programming infrastructure) is used to support communication between objects so that the reports may be from objects that are remote from the accounting object. The Accounting class implements the ReportListener interface so that it can receive those accounting reports as well. The accounting report events are of a

ReportEvent class. An event in this class is a traditional Java event which includes a Report object. The main attribute in the Report class is records. Records is a simple vector of accounting records. These records are described in the AccountingStoreView. The ActReportCtrl interface defines the control calls related to the accounting reporting process of an accounting object. Calls are available for an object to register interest in accounting reports, de-register interest and to control the reporting process.

The accounting store view (Figure 5c) regroupes the class related to the persistent storage of accounting information. An accounting object has a Database of accounting Records. The Record type holds accounting information which is not priced. Priced information is the subject of a different class. The Database class is a simple Vector of Record objects and it can be serialized to a file on an external storage medium. The database is also responsible for returning accounting records that have to be reported.

The accounting meter view (Figure 5d) regroupes the class related to metering aspect of the accounting class. This relates both to the reception of the measurement information in the accounting objects and also to the control of the Meter as well as the definition of a simple Meter class. The Meter class uses a "Pulsar" object that generates pulses events as required. The frequency of pulses is set by the Meter object. On reception of pulses the Meter generates objects of type MeasurementEvent. Objects implementing the MeasurementListener interface and that have registered their interest in measurement results will then receive those events via a measurementHandler method. As previously noted, the Meter object and one or more of the objects receiving measurement events may be remote from each other. A measurement event is a conventional Java event and includes a measurement record of type MeasurementRecord. An accounting object gets measurement information from a Meter over which it has got control via the MeterControl interface. A typical example of control is the measurement reporting frequency, that is, an accounting object may control the frequency with which a meter object sends reports to it. This control interface is also the one to use to register interest in measurement results.

The accounting miscellaneous view (Figure 5e) regroupes all the other classes that do not fit in the previously described views. This includes, JavaBean-

related classes, classes to run the code and graphical user interfaces (GUI). The AccountingBeanInfo class is a JavaBean related class which modifies the description of some attributes on the Accounting class when those properties have to be graphically displayed in the BeanBox or in any other component builder tool.

5 The Go and MeterGo classes only implement a main method. Go is used to launch an accounting object and MeterGo a Meter object. The AccountingGUI class is responsible for the GUI related to the accounting objects. The Meter object has no GUI associated with it. The Accounting GUI is shown in Figure 6. The top part of the GUI includes data from the Accounting object and the bottom part relates to

10 the control available over the accounting object. The control part is directly related to the control interfaces available for the Accounting objects. The accounting class is not aware of the GUI as the reference is from the GUI to the accounting class.

The accounting mechanisms described above can be used in combination with contracts between customers and retail and wholesale networks to establish

15 liability to pay and who is expected to pay. The following section describes different clearing models for the making of payments. The systems described in this section may be used in conjunction with, or independently of the specific accounting mechanisms described above.

Payment Clearing

20 As well as "liability to pay" and "who is expected to pay" there is also the question of who should be paid. It is preferable for it to be customary for each edge ISP to be paid on a "half-circuit" basis for both their sent and received service. However other business models need to be considered. In particular, we will now consider a model similar to the public phone service, which has one or

25 two implicit features that need to be separated out for full understanding.

Let us consider a business model where ISPs don't expect payment for all sent and received traffic to be made to all edge providers. Instead a customer might pay their own provider on behalf of both (all) ends as in telephony. A further accounting field would appear to be necessary - a "payee" field. For instance, this

30 alternative business model might be that the decision as to which end(s) payment from edge customers entered the system was made on a per flow basis by customers. We shall call this model the "provider clearing" model for reasons that will become clear as we go. This is shown in Figure 9. Here, end customers 91,92 communicate via a number of intermediate networks 93. The financial

flows between providers in this model depend on at which ends payment is entering the system on a per flow (or per packet) basis. For some flows, there may even be proportional sharing of costs between the ends. Therefore, for business model flexibility, rather than stating simply "local" or "remote" end, the "payee" field could be a "payee percentage" field instead - the percentage of the total cost to be paid by the customer at the end being accounted for. So usually it would be 100% or 0% in the typical cases of "paid completely to local provider" or "completely to remote". The balance would be the remote end's payment. Note, though, that the perceived purpose of this model is the transaction efficiency when the local payee gets 100%. However, there are certain disadvantages for the "provider clearing" model:

As already pointed out, the "payee percentage" field would have to drive inter-provider accounting, otherwise the revenue of an edge ISP and its upstream providers would depend on a factor completely outside their control - to which end its customers chose to make payment. The "payee percentage" field would therefore have to be trusted by upstream providers. To help prevent the field being tampered with, it would need to be signed by the remote ISP. How signed fields can be aggregated without losing the signature integrity is a matter for further research. The aggregation might have to be done by software signed by a third party trusted by all the parties involved (TTP) and then the record re-signed by the TTP. However the aggregation software would also have to run on a host trusted by the TTP. Further, using this model would mean that all edge ISPs would have to be able to identify any remote ISP from the remote address, something not possible with hierarchical routing. Nonetheless, we have already stated that the payment interface of the remote ISP can be passed in a higher level protocol between end stations. It would be only slightly more complex for them to include this in the accounting record. However, the ISP would still have to make appropriate checks that this was a valid ISP and that it matched the remote address. Once it has the address this becomes trivial, but more inefficient and rather negates the advantage of the local ISP doing the clearing via its upstream provider. Still further complication might be introduced for some future applications if the share of payment between the parties wasn't fixed but depended on characteristics of the flow or other parameters only understood at a higher level -

higher than the provider would normally be interested in. This is also a problem for the "expected payer" field, but in that case the field is informational only, unlike the "payee percentage" field in the "provider clearing" model.

Worse still, the payment should ideally be split taking into account the current
5 prices of all the edge providers who will eventually be paid. The only alternative (used in the international accounting rate system (IARS) for telephony) is for ISPs to agree compromise prices between themselves that average out price inconsistencies. This is what has been causing all the tensions in IARS as some countries liberalise earlier than others causing huge variation in prices around the
10 world, between which no compromise can be found with which all involved are content. This is difficult even for a system where every end to end path only passes through two international carriers at maximum, each pair setting compromise prices with each other. With nine ISPs on many end to end Internet paths and considerable peer interconnection, the horse trading would be a
15 nightmare.

Finally, because of the much longer provider chains typically found on the Internet, potentially unacceptable delays will be introduced before the revenue arrives in the correct place. Any delay in clearing hugely increases the cost of the payment system, as extra trust mechanisms have to be invoked while the payment
20 remains unconfirmed. These trust mechanisms have to be applied to the edge customers, not just the providers, therefore hugely increasing the total cost of the system.

Despite this limitations, the reason such a model is appealing is that it appears to reduce the number of payment transactions. For example, if the parties in an
25 Internet 'phone conversation are both (all) being paid for by the caller, it appears less complex for the caller to pay everyone's payments to her own ISP, then let the ISP transfer the correct amount to its upstream provider as part of a bulk transaction. On the other hand, in a "third party clearing" model (shown in Figure 10), the caller has to split up the payment between both (all) ISPs of both
30 (all) parties involved.

This is why the distinction between the names of the two models is in the clearing, not who is paid. Both models end up with edge ISPs paid on a half-circuit basis. The difference is merely in the route the payment takes from payer to payee. With

provider clearing the payment follows the data path. Along the way, providers take their cut with two types of money sharing being mixed together: wholesale cut half-circuit sharing

In the "third party clearing" model, the clearing house rôle deals with the half-circuit sharing (including the straightforward price differences between the two ends) leaving inter-provider accounting to be purely about wholesaling. There is nothing to stop providers or customers assuming the clearing house rôle, but the accounting information model needs to be based on a third party clearing system to allow for the most general case. To clarify, whether the paying customer makes payment to a dedicated clearing house, direct to the ISP at the remote end or even direct to the remote customer so that they can pay their own ISP, in all cases, the rôle of clearing must be separate even if there is no separate enterprise to achieve the function. Note that the last case is special - the clearing rôle is null, but it still appears in the information model. In other words, the charges for all ends should never be lumped together while accounting. If the half-circuit sharing is achieved through the provider chain, this must be kept separate from the accounting for wholesale. If it is not, the types of model that can be built on the infrastructure are restricted.

Having separated out the rôle of clearing, this now shows explicitly that a telephone company also bundled another rôle in its business- that of "session retailer". That is, the edge telco is offering telephony sessions at fixed prices, but the range of prices is less than the number of possible ways the price could vary if it were simply composed of all the end to end prices charged by providers necessary to assemble each session. Again, this rôle may be assumed by the edge customer in the Internet world, but it is possible that businesses will spring up offering prices for transmissions by selling on IP service while absorbing variations across providers in the prices they are charged wholesale. Obviously, this rôle may also continue to be taken by telcos and ISPs too.

It is redundant to state in accounting messages which end will actually be paid. Who should eventually receive the payment is implicit because the rule is now that accounts for other providers shouldn't be lumped with accounts for the local provider. The corollary is that any accounting implicitly relates to payments that will eventually end up with the local provider. Saying who will be paid is meaningless during accounting. It is only relevant at the time of payment. Then it

essential to say who the payment is eventually intended for if it is given to a clearing organisation.

Other aspects of the invention are described, by way of example, in the annexed paper.

The Direction of Value Flow in Connectionless Networks

Abstract

This paper proposes solutions to an apparently fundamental question in networking: "In which direction does value flow in a connectionless communications network?" Value flow is considered both between the ends of a communication and between the networks along the path of a communication. Multicast and aggregation modes are considered as well as unicast. The goal is to derive an optimal default business model for Internet service providers if charging for transmission. The search is for the most likely common case for apportioning the value of transmission between senders and receivers, in order to reduce the cost of dealing with exceptions to the norm. But this has to be reconciled with conflicting issues of blame, liability and control. A pre-requisite is to unravel the confusions that are common where higher level issues become embroiled with networking issues. The result is the creation of two possibly novel business models.

Keywords

Charging, pricing, end to end, Internet.

1. Introduction

The value of communication concerns the value of having information in a certain place or places, rather than the value of the information itself. However, usually, the more the information is worth, the more value is placed on having it in the right place. But, because the data communications market is fairly competitive, charges for communicating information tend instead to follow the cost plus margin rule. This is particularly so because there is no way for providers to know what value their customers put on moving any one piece of information anyway.

Traditionally, data communications has been sold so cheaply that charging for it on usage basis has not seemed feasible or sensible. Where flat-rate subscription prevails, the question of the apportionment of the value of a particular communication between its ends rarely surfaces. With the possibility of variable quality of service (QoS) approaching, the need for some form of usage-charging for high QoS service has arisen. This has led to new thinking on cheaper usage-charging systems for packet networks [xxx99]. This brings the issue of the apportionment of the value of a communication back into the limelight.

Any payment to an edge-network provider has the two aspects - 'who pays' and 'who is paid'. 'Who is paid' can only be each local provider collecting its local price. With 'cost plus' pricing there is no scope for any provider to break out of that. But, because communications naturally involves at least two parties, there is a clear opportunity that, in order to cover the total costs of all the providers involved, 'who pays' can be on a different apportionment.

The edge customers *do* know the value to them of having the information at a certain place. Thus, although apportionment is irrelevant to the network providers (as long as they get paid), it is very relevant to the edge-customers. Also, clearly, the network providers can stimulate more use of their networks by making arrangements for customers to efficiently apportion costs to their whim.

Clearly, the cost of apportioning charges is significant, therefore it is important that the default apportionment matches the most common case. This paper establishes that default case then goes on

apportionment matches the most common case. This paper establishes that default case then goes on to suggest a new role in the communications industry for apportioning charges between end-customers separately from the question of the apportionment of payment to network providers. This is contrasted with the current way apportionment is universally achieved in the communications industry and with some of the proposals in the literature.

The asymmetric nature of the relationship between sender and receiver is also discussed, with respect to blame, liability and control over the flow of information.

2. Related Work

Some of the literature is very sloppy in respect of the direction of value flow, even when only considering fixed Internet access charges. Some authors even state that they believe the current Internet model is "sender takes all" [ITU96, Zull97]. That is, they assert that the only revenue received for access bandwidth is by the sender's ISP, as if traffic out of a network doesn't need any bandwidth.

MacKie-Mason *et al* asserts that blame is impossible to determine at the network level [MacKieVar92], an argument that can descend into sophistry, as both concepts are difficult to define. Clark analyses the apportionment of charges between senders and receivers [Clark96], and proposes an engineering solution, which it is admitted would introduce considerable complexity to the Internet if implemented. Shenker *et al* describes edge pricing [Shenker96], a business model that is prevalent in communication networks and which forms much of the background to this work.

3. The direction of traffic value flow

In general, we assert that the value of the networking service flows from the network provider outwards to each of its customers, whichever direction traffic is flowing. This is because the large majority of transmissions are with the consent of all ends. If operating usage-based charging, we propose a provider should aim to offer each type of packet transmission service in each direction at a separate price. There is obviously nothing to stop any two prices being the same. Types of service are defined both by their service mode (unicast, multicast) and their quality (latency, instantaneous bandwidth, reliability, jitter).

If a price is higher than the perceived value for any customer, she is free to get the remote party (or anyone else) to make up the difference through some higher level arrangement. On the other hand, if the value to her is higher than her local price, she is also free to offer to cover some of the costs of the remote end(s). However, the provider in our minimalist business model shouldn't be concerned with matchmaking multiple customers to get round local discrepancies between price and customer value. This is an issue that should be dealt with end-to-end not locally. We are not saying ISPs shouldn't offer end-to-end pricing - it is clearly in their interest to matchmake between customers with surplus value and those with deficit. All we are saying is that, if they do, end-to-end pricing should be considered as a separate role (Fig {3.1?}). Such a role could be a separate business - it could gain on some combinations and lose on others, possibly making a profit overall. It would be a retail service that uses the networking services as wholesalers. It is also possible that edge customers could effectively take on this role themselves. Fig {3.1?} shows two end customers connected by multiple ISPs. The relative value of the service flows and prices for one direction of one class of service is represented by the thickness of the arrows. Note that the size of the proportions of prices represents a choice by the end-system willing to pay more than its local price. Pricing between providers is omitted for clarity (but see later).

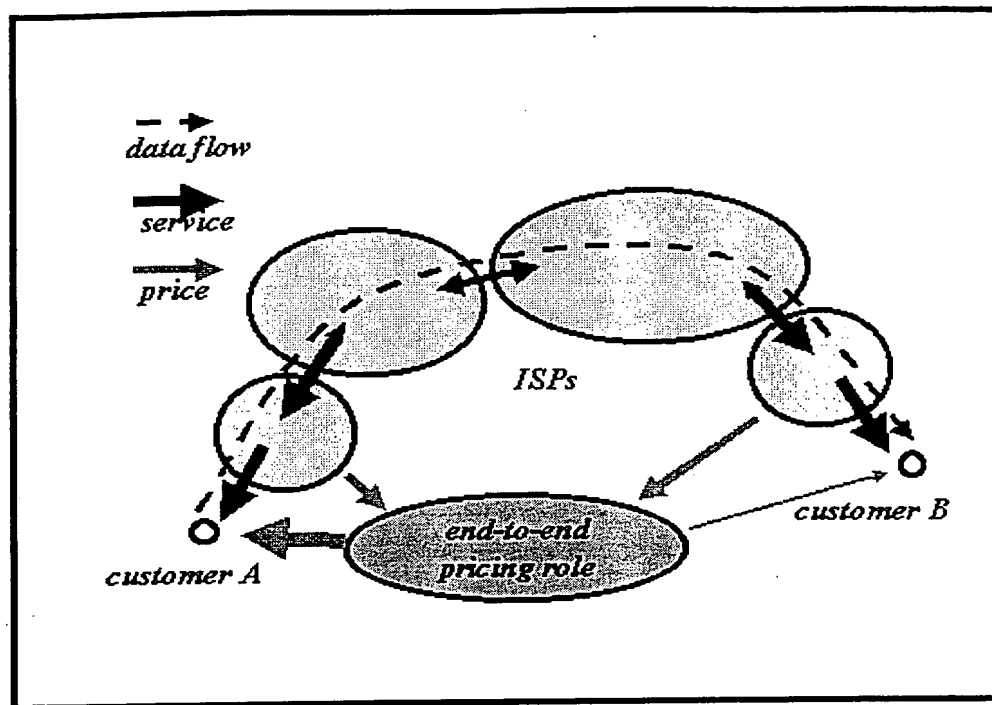


Fig {3.1?} - End-to-end pricing role

Telephony firms have traditionally offered end-to-end pricing because they are selling an application. The role of network provider has always been muddled with selling the end-to-end application. This is already putting considerable strains on the International Accounting Rate System (IARS) [ITU RIARS] with potentially $s(n-1)^2$ prices having to be negotiated (where n is the number of edge providers and s is the number of global schemes for sharing the proportions of the price between the ends, e.g. local rate only, free to sender). In practice, end providers are grouped together to reduce the number of prices presented to customers. The PSTN uses addressing conventions (e.g. +800 for free to sender), but this limits commercial flexibility to the few schemes that are widely recognised. Clark proposed a solution to allow flexibility [Clark96]. However, catering for various combinations of sender and receiver payments through the core of the network needs packet format changes and router involvement. Further, wholesale prices between providers would have to be negotiated for every possible scheme for sharing charges between the two ends as well as for every possible grouping of end points beyond that boundary. Worse still, inter-provider accounting would then require traffic flows to be isolated then further sub-classified by how much each end was paying on a per-packet basis.

The ' n^2 problem' would still exist for our end-to-end pricing solution but this is fairly easy to contain by grouping. Importantly though, end-to-end pricing gets rid of all the inter-provider problems described above. There becomes no need at all to identify end-to-end flows at inter-provider boundaries. Thus inter-provider charging could be based on bulk measures like average queue lengths, routing advertisements etc. Also, most importantly, end-to-end pricing can be introduced without changing the Internet at all, and it allows future flexibility. To summarise so far, we should ensure any discrepancy in the willingness to pay across the end customers is normalised end-to-end first, so that edge ISPs always receive payment at their local price.

Although we have delegated the problem of combining sender and receiver payments to a higher level, we still believe it is important to cater for the common case at the network charging level so that the higher level functions are unnecessary in most cases. The large majority of communication occurs between consenting parties. This is why we proposed above that all edge providers should charge their local customers for both sending and receiving. Allowing different prices for each direction allows for asymmetric costs (e.g. access technology like ADSL or satellite) and for asymmetric demand (e.g. some ISPs might host more big senders, while others might host the mass of receivers). Now that we have eliminated all but local pricing, we can extend this model recursively to apply at the boundary between any pair of providers. This becomes a generalised

model of edge pricing whether the 'edge' is really at the network edge or just the edge of a backbone.

Fig {3.2?} shows a generic scenario to analyse whether any one choice between sender and receiver payments is more stable. It shows multiple networks, N_i , all connected to the network of interest, N_b . For each class of service (type of packet), each connected network has a status relative to N_b based on whether it provides more or less connectivity to other hosts at that class of service. Although the diagram gives the impression that N_b is a backbone network, any one of the neighbouring networks could be a simple link to an edge customer's single host. The model is designed to be general enough for N_b to be an edge customer, an edge network, a backbone network or some hybrid. Those networks with the same suffix are of similar status relative to N_b . For instance, those labelled N_c may be edge customers, N_d may be equally large backbones and N_e a peer network.

A packet is shown being multicast from N_a into N_b and onward into the other networks. Because multicast is a general case of unicast this allows us to model both topologies. We will also be able to treat the topology as aggregation⁽¹⁾ by reversing the direction of transmission. The term packet is used, but the arrows could represent flows of similar class packets for a certain time. The packet or flow being modelled could be data or signalling. It is not necessary to model multi-source multicast separately because packets from different sources always remain separate. Fig {3.2?} highlights the pricing between networks N_a and N_b . W_{bas} and W_{bar} denote the per direction weightings applied to the charge that N_b applies to N_a . W_{abs} and W_{abr} likewise weight the charge N_a applies to N_b . Each weighted price is for transmission between the edge in question and the remote edge of the Internet, not just the remote edge of that provider. There would be four price weightings like this at every inter-network interface, but the weights would take different values unless the neighbours were of the same status. Thus the payment for traffic in any one direction across each interface depends on the difference between the two weighted prices offered by the networks either side. In other words, no assumptions are made about who is provider and who is customer; this purely depends on the sign of the difference between the charges at any one time. Clearly, edge customers (N_c , say) have no provider status in the networking market. So, for all j , $W_{cjs} = 0$ and $W_{cjr} = 0$. We can then analyse scenarios like 'only senders pay' or 'only receivers pay' by setting all receiving weights to zero or all sending weights to zero. For instance, stability of a policy can be determined by assessing whether one network would gain from a maverick policy.




Fig {3.2?} - Charge for sending, receiving or both?

'Only senders pay' or 'only receivers pay' tends to encourage migration of customers who are

primarily receivers and those who are primarily senders to different providers. This situation is tenable because the provider with all the non-paying customers gets all its revenue from its interconnect business. Either scenario remains stable, because if one network goes maverick (e.g. only charges receivers when everyone else is only charging senders), both predominant senders and receivers have a choice of cheaper provider. Therefore the income to the whole system reduces ensuring the maverick provider would go bust first - sufficient disincentive to be maverick! However, both these policies clearly make network utilisation inefficient and both are unstable where multicast (and consequently aggregation) are concerned.

In contrast, 'senders and receivers pay' is stable in both unicast and multicast cases. It also doesn't lead to inefficient network utilisation unlike the above cases. It is also possible to cater for different balances of predominant senders and receivers by weighting the sending price differently to the receiving price. For instance if there are a few big predominant senders but many small predominant receivers, the economy of scale in managing a large customer can be reflected in a lower sender weighting. Similarly, the inefficiencies of multicasts to small receiver communities compared to multiple unicasts can be discouraged by slightly weighting multicast sender pricing.

We have shown that all ends paying is both the common case and a stable one so should be the default. We can share the cost differently at higher level if end user value is shared differently from this default (and it is worth bothering given the cost of another financial transfer). However, if it is the receiver that has the deficit in value, we must remember that, in the final analysis, a sender can decide not to send but a receiver can't avoid being sent to (in the current Internet). Such cases are much rarer than it first appears, mainly because of confusions that can be cleared by considering the following factors:

- The value of the information isn't relevant when considering the networking service - only the value of *moving* the information - getting it to a useful place
- Often the value of moving information is transitory - getting it to a useful place to discover that moving it wasn't useful
- Often the value of moving lots of information is to get a small part of it to a useful place, but it isn't possible to know which part before moving it
- The cost of transmitting information is often far less than the cost of targeting which information should be transmitted
- Information in one direction often controls the flow of information in the other

Nonetheless, genuine cases remain where the receiver is being persistently forced to pay for transmission that is valuable to the sender but not to the receiver. The only solution to this seemingly intractable dilemma is for it to be *customary* for all ends to pay, but the ultimate liability should remain with the sender. Any receiver could then dispute the customary apportionment (end-to-end) with no risk of denial (unless the sender had proof of a receiver request). A similar but opposite situation used to prevail with the UK postal service. It was customary for the sender to pay for the stamp, but if it was missing or insufficient the receiver was liable for the payment, because the Royal Mail had an obligation to deliver every letter.

4. Clearing

We are assuming electronic commerce will make it possible for anyone to pay anyone else's ISP on the Internet, even if a clearinghouse is needed. These arrangements will be made through higher level (typically session level) protocols. It is assumed that if one customer wants her ISP to be paid by a remote customer, she will communicate her ISPs payment interface address as part of the high level protocol. We shall call this the "third party clearing" model. This fits our earlier assertion that it would be sensible for it to be customary for each edge ISP to be paid on a "half-circuit" basis for both their sent and received service. However other business models need to be considered. In particular, we will now consider a model similar to the public phone service, which has one or two implicit features that need to be separated out for full understanding.

Let us consider a business model where ISPs don't expect payment for all sent and received traffic to be made to *all* edge providers (Fig {4.1?}). Instead a customer might pay their own provider on behalf of both (all) ends as in telephony and [Clark96]. For instance, this alternative business model might be that the decision as to which end(s) payment from edge customers entered the system was

made on a per flow basis by customers. We shall call this model the "provider clearing" model for reasons that will become clear as we go. The financial flows between providers in this model depend on at which ends payment is entering the system on a per flow (or per packet) basis. For some flows, there may even be proportional sharing of costs between the ends. For business model flexibility an accounting system would need a "payee percentage" field - the percentage of the total cost to be paid by the customer at the end being accounted for. Usually it would be 100% or 0% in the typical cases of "paid completely to local provider" or "completely to remote". The balance would be the remote end's payment. Note, though, that the perceived purpose of this model is the transaction efficiency when the local payee gets 100%.




Fig {4.1?} - "Provider clearing" model

However, there are five points stacked up against the "provider clearing" model:

- As already pointed out, the "payee percentage" field would have to drive inter-provider accounting, otherwise the revenue of an edge ISP and its upstream providers would depend on a factor completely outside their control - to which end its customers chose to make payment. The "payee percentage" field would therefore have to be trusted by upstream providers. To help prevent the field being tampered with, it would need to be signed by the remote ISP. How signed fields can be aggregated without losing the signature integrity is a matter for further research.
- Further, using this model would mean that all edge ISPs would have to be able to identify any remote ISP from the remote address. Nonetheless, the payment interface of the remote ISP can be passed in a higher level protocol between end stations. It would be only slightly more complex for them to include this in the accounting record. However, the ISP would still have to make appropriate checks that this was a valid ISP and that it matched the remote address. Once it has the address this becomes trivial, but more inefficient and rather negates the advantage of the local ISP doing the clearing via its upstream provider.
- Still further complication might be introduced for some future applications if the share of payment between the parties wasn't fixed but depended on characteristics of the flow or other parameters only understood at a higher level - higher than the provider would normally be interested in.
- Worse still, the payment should ideally be split taking into account the current prices of all the edge providers who will eventually be paid. The only alternative (used in the international accounting rate system (IARS) for telephony) is for ISPs to agree compromise prices between themselves that average out price inconsistencies. This is what has been causing all the tensions in IARS as some countries liberalise earlier than others causing huge variation in prices around the world, between which no compromise can be found with which all involved

- The content. This is difficult even for a system where every end to end path only passes through two international carriers at maximum, each pair setting compromise prices with each other. With eight ISPs on many end to end Internet paths, five typical [McCreary98] and considerable peer interconnection, the horse trading would be a nightmare.
- Finally, because of the much longer provider chains typically found on the Internet, unacceptable delays will be introduced before the revenue arrives in the correct place. Any delay in clearing hugely increases the cost of the payment system, as extra trust mechanisms have to be invoked while the payment remains unconfirmed. These trust mechanisms have to be applied to the edge customers, not just the providers, therefore hugely increasing the total cost of the system.

If so much about this business model is complicated, why is it even being discussed? The reason such a model is appealing is that it appears to reduce the number of payment transactions. For example, consider the case where both the parties in an Internet 'phone conversation are being paid for by the caller. It appears less complex for the caller to pay everyone's payments to her own ISP, then let the ISP transfer the correct amount to its upstream provider as part of a bulk transaction. On the other hand, in the "third party clearing" model, (Fig {4.2?}) the caller has to split up the payment between both ISPs of both parties involved.




Fig {4.2?} - "Third party clearing" model

This is why the distinction between the names of the two models is in the clearing, not who is paid. Both models end up with edge ISPs paid on a half-circuit basis. The difference is merely in the route the payment takes from payer to payee. With provider clearing the payment follows the data path. Along the way, providers take their cut with two types of money sharing being mixed together:

- wholesale cut
- half-circuit sharing

In the "third party clearing" model, the clearinghouse role deals with the half-circuit sharing (including the straightforward price differences between the two ends) leaving inter-provider accounting to be purely about wholesaling.

There is nothing to stop providers or customers assuming the clearinghouse role, but the accounting information model needs to be based on a third party clearing system to allow for the most general case. To clarify, whether the paying customer makes payment:

- to a dedicated clearing house

- direct to the ISP at the remote end
- or even direct to the remote customer so that they can pay their own ISP

in all cases, the *role* of clearing must be separate even if there is no separate enterprise to achieve the function. Note that the last case is special - the clearing role is null, but it still appears in the information model. In other words, the charges for all ends should never be lumped together while accounting. If the half-circuit sharing is achieved through the provider chain, this must be kept separate from the accounting for wholesale. If it is not, the types of model that can be built on the infrastructure are restricted.

5. Limitations and Further Work

TBA

6. Conclusions

We have shown that the common case for apportioning value between the ends of a connectionless communication network is catered for if all users pay for both sending and receiving. We have also shown that this is the most stable and efficient case, particularly for multicast and aggregation. It should therefore be the default apportionment for payment purposes.

We have suggested that a new business model would be useful and more efficient to cater for the cases where there is a large discrepancy from this default in terms of value apportionment - large enough for it to be worth making a balancing transaction given the costs of doing this. This new model requires a new role in communications markets - an end-to-end pricing role. In discussing clearing of payments across an end-to-end path, there is also a need for a third party role for end-to-end clearing. These two roles only make sense as new types of business if they are enacted by the same business. Otherwise customers will be paying money to a different organisation than the one quoting prices, which has obvious security flaws.

This new role could be conducted by existing ISPs or customers themselves, but there appears to be considerable added value, making this a viable business in its own right. It appears that this role is a threat to existing ISPs business. This role turns edge ISPs into wholesalers for a potentially large class of Internet applications. The end-to-end pricing and clearing role would become the retail face of the Internet in many cases.

Further, we suggest a subtle twist to the recommendation that customers should pay for both sending and receiving. We suggest this should be customary, but that ultimate liability for sending should lie with the sender. Disputes could then quickly be resolved through the end-to-end clearing role.

7. Acknowledgements

8. References

[McCreary98] Sean McCreary and kc claffy, "How far does the average packet travel on the Internet?", CAIDA, 25 May 1998, <URL:<http://www.caida.org/Learn/ASPL/>>

[Clark96] David D Clark (MIT), "Combining Sender and Receiver Payments in the Internet", in Interconnection and the Internet. Edited by G. Rosston and D. Waterman, Oct 1996, <URL:<http://diffserv.lcs.mit.edu/>> {not discussed in this paper yet - see .bib annotation}

[Finlayson98] Ross Finlayson, Discussion at Third Reliable multicast research group meeting, Orlando, FL, and following on the mailing list, Feb 1998 <URL:<http://www.east.isi.edu/rm/>>

[ITU96] ITU, "The Direction of Traffic", ITU/Telegraphy Inc, Geneva, 1996 <URL:<http://www.itu.ch/ti/publications/traffic/direct.htm>> in brief chapter 3, Box 3.1 is extracted on-line: "Accounting rates and how they work", <URL:<http://www.itu.ch/intset/whatare/howwork.html>>

[ITU 'IARS] ITU "Reforming the International Accounting Rate System"
<URL:<http://www.itu.ch/intset/>>

[MacKieVar92] Jeffrey K MacKie-Mason and Hal Varian, (UMich), "Some Economics of the Internet", Tenth Michigan Public Utility Conference at Western Michigan University, March 25-27, 1993: <URL:<http://www.sims.berkeley.edu/~hal/people/hal/papers.html>>

[Shenker96] Scott Shenker (Xerox PARC), David Clark (MIT), Deborah Estrin (USC/ISI) and Shai Herzog (USC/ISI), 'Pricing in Computer Networks: Reshaping the research agenda', SIGCOMM Computer Communication Review Volume 26, Number 2, Apr 1996, <URL:<http://www.statslab.cam.ac.uk/~frank/PRICE/scott.ps>>

[Speakman98] Tony Speakman, Dino Farinacci, Steven Lin, Alex Tweedly, (cisco) "PGM Reliable Transport Protocol Specification", Work in progress: IETF Internet Draft, Jan 1998 (expires Jul 1998), <URL:draft-speakman-pgm-spec-01.txt>

[xxxb99] XXX, "Lightweight, End to End, Usage-based Charging for Packet Networks", Submitted to ACM SIGCOMM'99, Jan 1999, <URL:yyy> (attached)

[Zhang93] Lixia Zhang (Xerox PARC), Stephen Deering (Xerox PARC), Deborah Estrin(USC/ISI), Scott Shenker (Xerox PARC) and Daniel Zappala (USC/CSD), "RSVP: A New Resource ReSerVation Protocol", IEEE Network. Sep 1993. <URL:<http://www.isi.edu/div7/rsvp/pub.html>>

[Zull97] Chris Zull (Cutler & Co), "Interconnection Issues in the Multimedia Environment", Interconnection Asia '97, IIR Conferences, Singapore, Apr 1997
<URL:<http://www.cutlerco.com.au/core/content/speeches/Interconnection%20Issues/Interconnection>>

Notes

(i) Examples of packets that are forwarded until aggregation (reverse multicast) are:

- RSVP[Zhang93] receiver initiated reservation (RESV) messages
- pragmatic general multicast (PGM) [Speakman98] negative acknowledge (NACK) messages or the "lay breadcrumb" messages[Finlayson98] suggested in their place

CLAIMS

1. A method of operating a communications network comprising:
 - a) measuring at each of a plurality of customer terminals usage by the
5 respective customer terminal of network resources; and
 - b) subsequently calculating a network usage charge from the
 measurement data generated by step (a).
2. A method of operating a federated data communications network characterised
10 by measuring at each of a plurality of customer terminals connected to the said
 network usage by the respective customer terminal of network resources.
3. A method according to claim 2, further comprising subsequently calculating a
 network usage charge from measurement data generated by the step of
15 measuring.
4. A method according to any one of the preceding claims, further comprising
 step of aggregating measurement data produced by a series of measurements
 at respective customer terminal.
20
5. A method according to any one of the preceding claims, further comprising
 storing the measurement data.
6. A method according to claim 5, including storing with the measurement data
25 data identifying a tariff applicable to the said measurement data.
7. A method according to any one of the preceding claims including
 communicating data generated by step (a) to a network accounting object
 controlled by a network operator.
30
8. A method according to claim 7, including communicating to the network
 accounting object a usage charge calculated from the measurement data.

9. A method according to any one of the preceding claims, including communicating measurement data to a system remote from the customer terminal.
10. A method according to any one of the preceding claims, including a step
5 carried out by the network operator of sampling part only of the traffic communicated between a customer terminal and the network and for the sampled traffic comparing the network usage with data communicated from the customer terminal to the network accounting object and thereby detecting any discrepancy.
- 10 11. A method according to any one of the preceding claims in which a network accounting object is configurable to receive data either from a measurement object controlled by the network operator or from a customer terminal.
12. A method according to claim 11, in which a customer accounting object
15 associated with the customer terminal is configurable to direct data to the network accounting object.
13. A method according to claim 11 or 12, including switching the network accounting object from a first configuration in which data is received from the said
20 measurement object and another configuration in which data is received from the customer terminal in response to a control signal received at the network accounting object.
14. A method according to any one of the preceding claims further comprising
25 communicating a tariff to each of the customer terminals, and calculating at each of the terminals from the tariff and from the accounting data the network usage charge.
15. A method according to any one of the preceding claims in which the
30 communications network is a federated data network comprising a plurality of network domains.
16. A method according to claim 15 including

communicating traffic between a customer terminal and a first network domain connected to the customer terminal,

further communicating the said traffic between the first network domain and a second network domain connected to the first network domain;

5 communicating network usage data from the customer terminal to a first network accounting object in the first domain;

communicating accounting data between the first network accounting object and a second network accounting object in the second domain.

10

17. A method according to claim 16, including determining from a current routing table in the first network domain the identity of a second domain, which second domain is communicating data with the customer terminal via the first network domain, and communicating network usage data for the customer terminal to the
15 second domain identified by the current routing table.

18. A method according to any one of the preceding claims in which the step of measuring includes counting the number of packets communicated between the customer terminal and the communications network.

20

19. A method according to claim 18, including counting both the number of packets received by the customer terminal and the number of packets sent by the customer terminal.

25 20. A method according to any one of the preceding claims, in which a payment for network usage is made to a third-party clearer.

21. A communications network arranged to operate by a method according to anyone of the preceding claims.

30


22. A customer terminal arranged to operate by a method according to any one of the preceding claims.

23. A customer terminal including a data interface arranged to be connected to a federated data network, characterised by a network usage meter arranged to measure the usage by the customer terminal of network resources.
- 5 24. A customer terminal according to claim 23, in which the usage meter includes means for counting the number of packets communicated between the customer terminal and the network via the data interface.
- 10 25. A customer terminal according to any of claims 22 to 24, including an accounting interface arranged to communicate measurement data to a network accounting object.
- 15 26. A method of operating a network comprising a plurality of network domains, including calculating a charge for use by a respective customer of network resources, and making payment in settlement of the said charge to a third party clearer.
- 20 27. A method according to any one of claims 1 to 20, including automatically varying a tariff for network usage in dependence on loading of the network, and calculating a charge for network usage by applying the tariff to the measurement data.

ABSTRACT

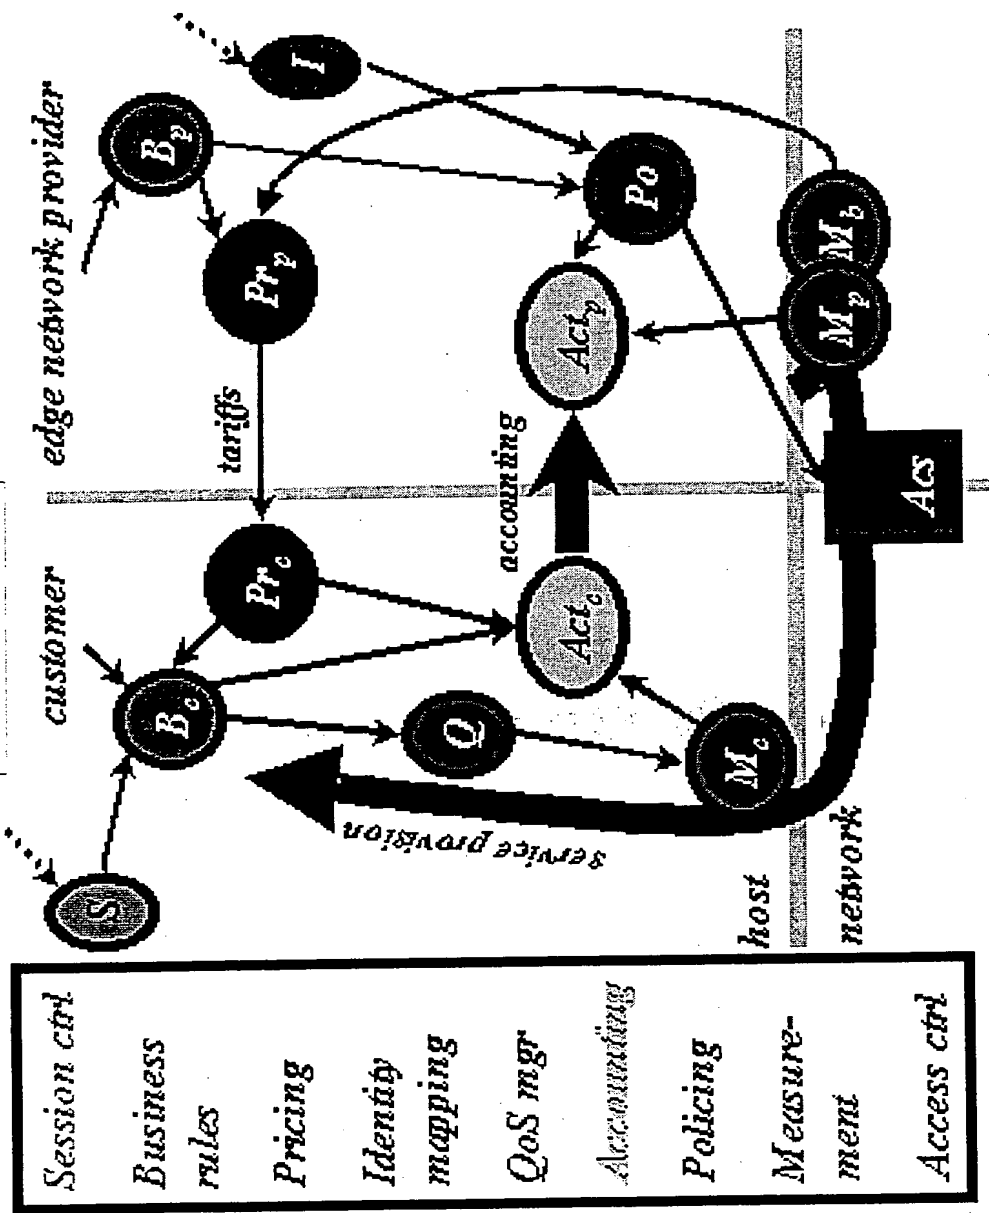
In a communications network, which may be a federated data network such as the
5 Internet, use of network resources is measured locally at customer terminals, for
example by counting the number of packets sent and received. The resulting data
may be aggregated and sent to a network accounting object. Accounting data
may subsequently be passed between network subdomains .

Figure 1



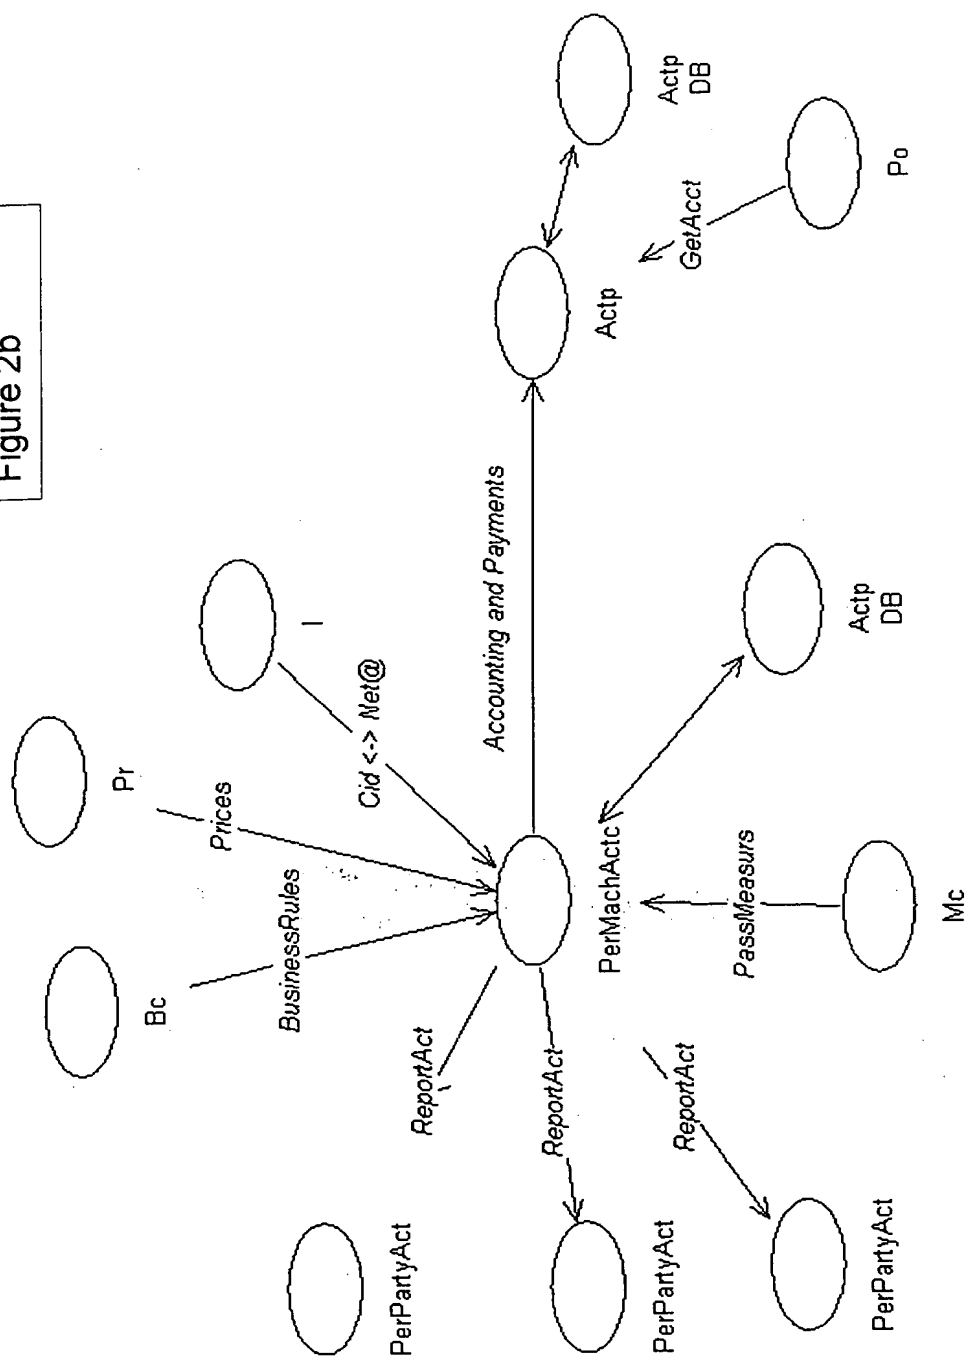
THIS PAGE BLANK (USPTO)

Figure 2a



THIS PAGE BLANK (USPTO)

Figure 2b



THIS PAGE BLANK (USPTO)




Figure 3a

accnte, tx_id, svc_id, unites, src, dest, time, period, tariff_id, cost, cur





Figure 3b

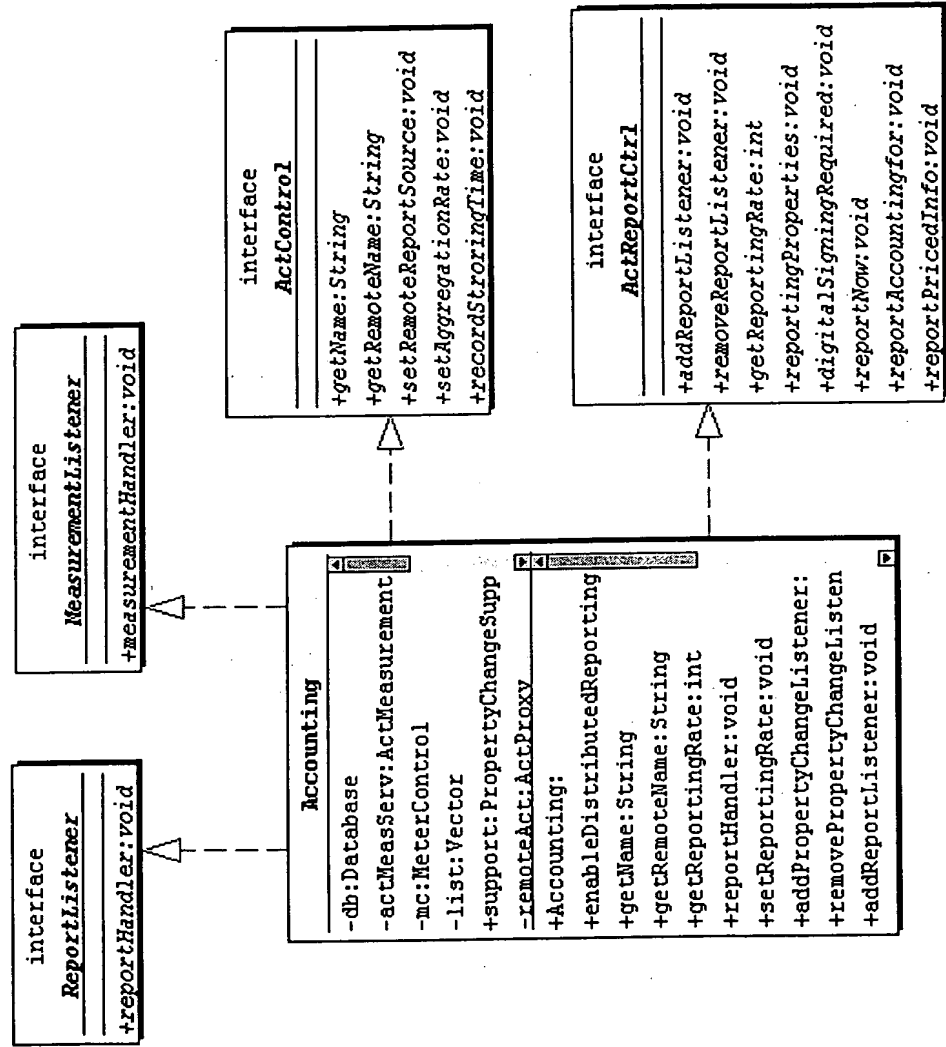
THIS PAGE BLANK (USPTO)

Figure 4



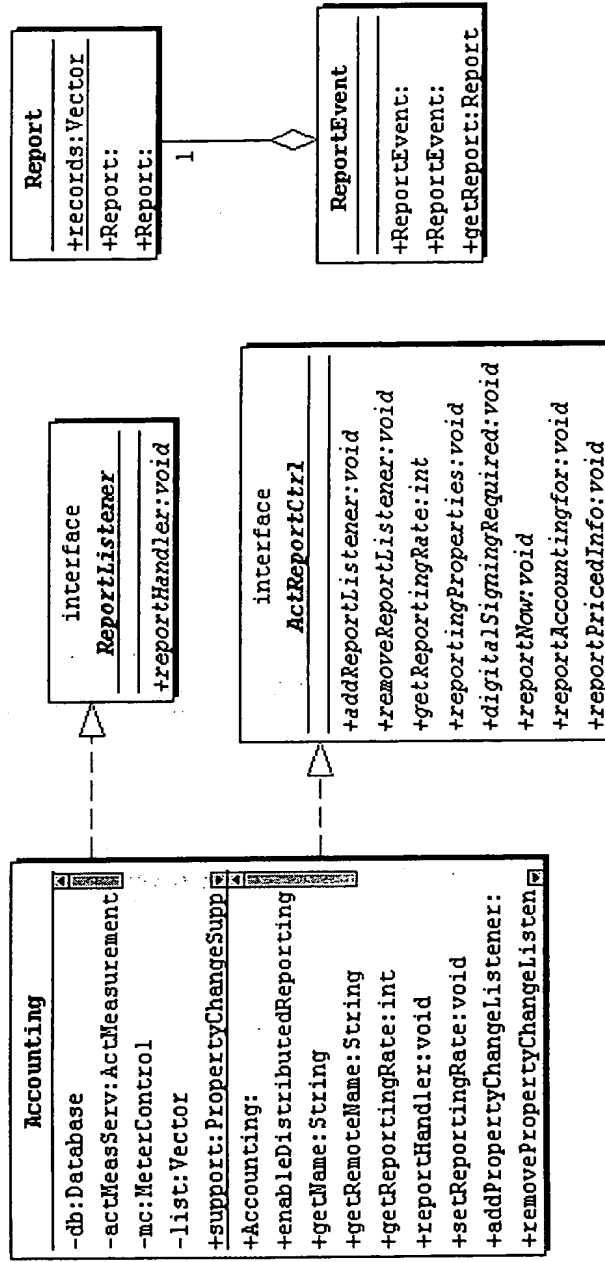
THIS PAGE BLANK (USPTO)

Figure 5a



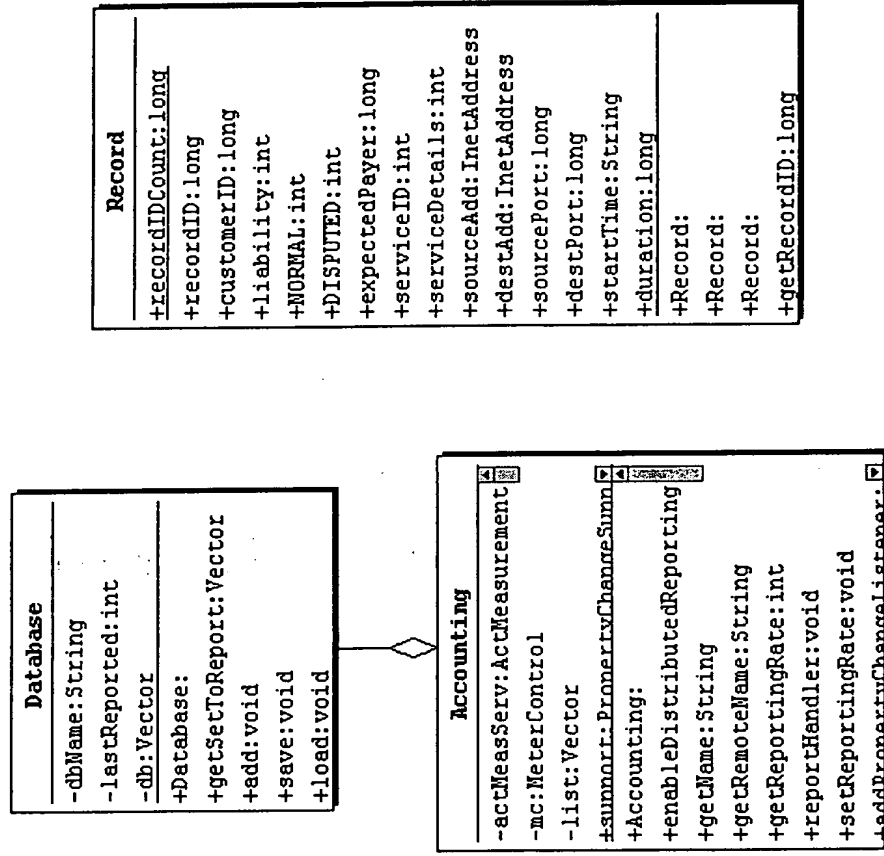
THIS PAGE BLANK (USPTO)

Figure 5b



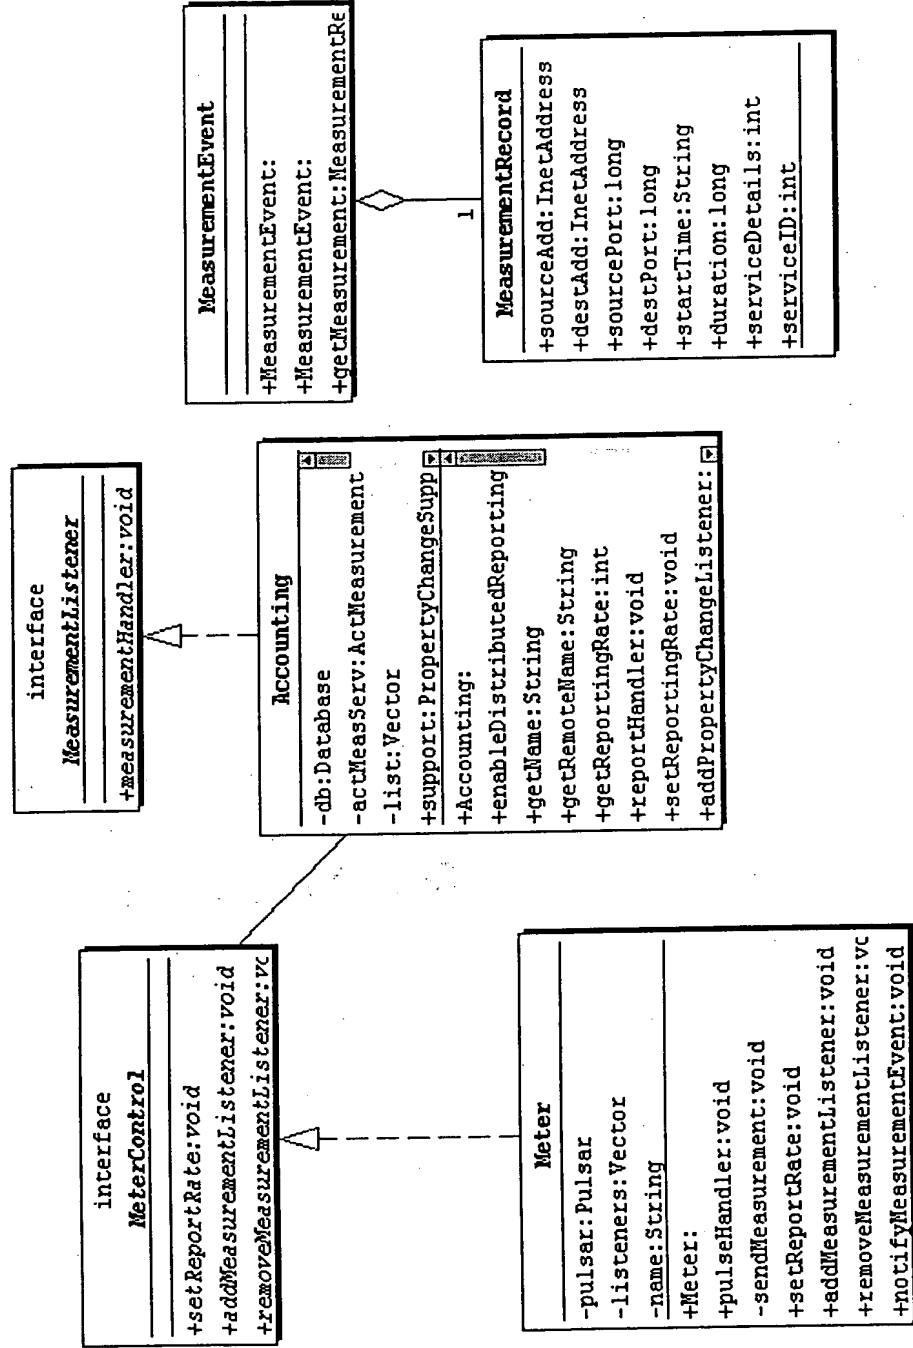
THIS PAGE BLANK (USPTO)

Figure 5c



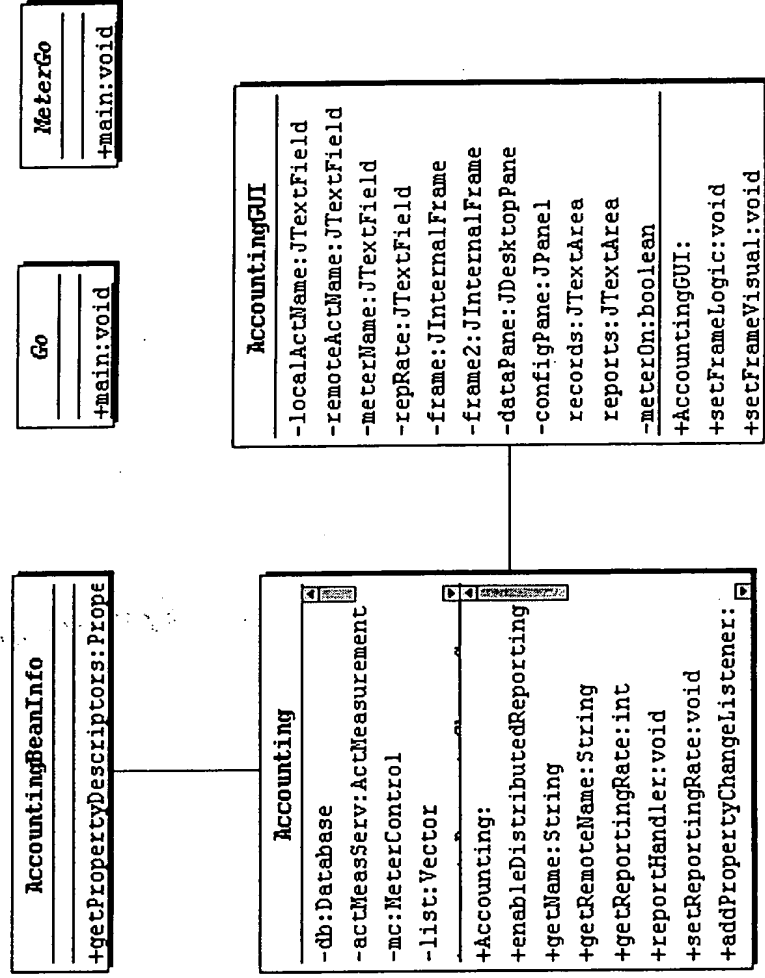
THIS PAGE BLANK (USPTO)

Figure 5d



THIS PAGE BLANK (USPTO)

Figure 5e



THIS PAGE BLANK (USPTO)

Figure 6

BTInternet Internet Accounting Control Platform

Local Platform ID

BTInternet

Local Meter ID

local

Local Reporting Rate

1000

Reporting Source

Add

Remove

Demon

MCI

BTInternet

Report NOW

Required Reporting Rate

2000

Reporting Phase

☐ Report priced data

Security

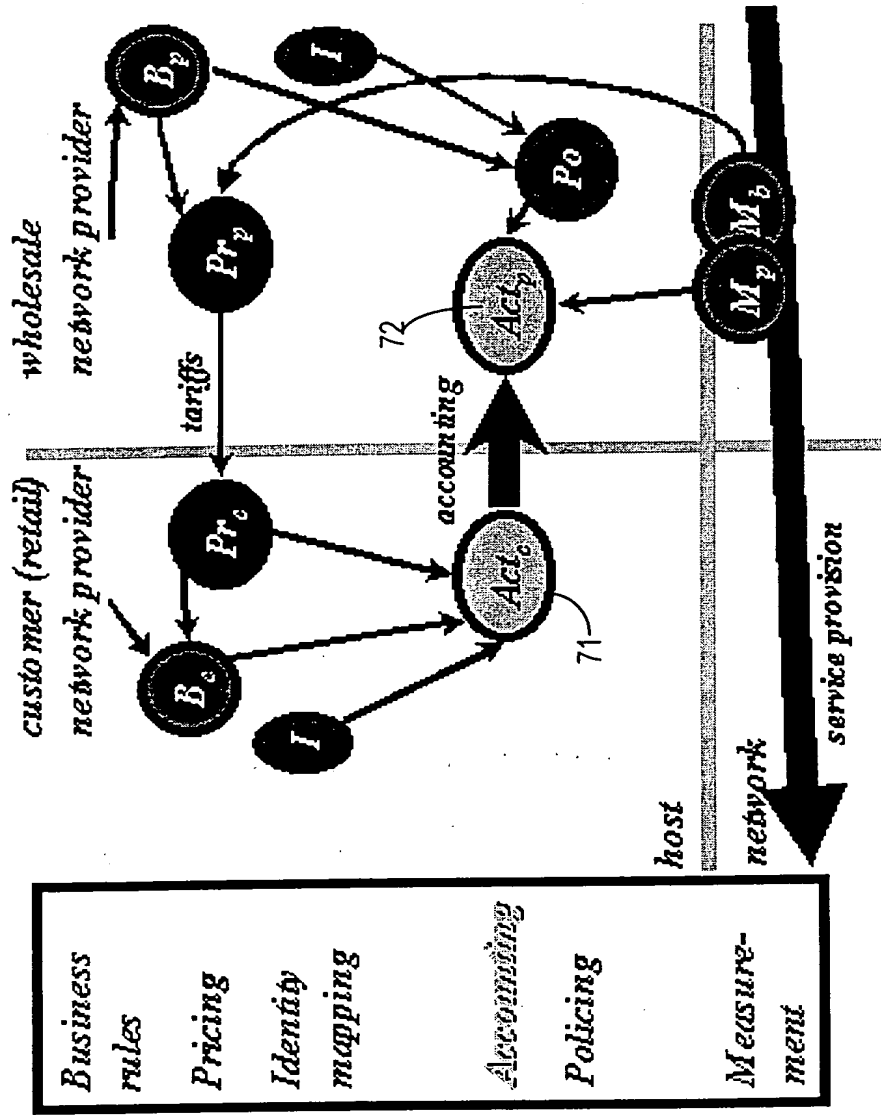
☐ Encryption required

☐ Digital signing Required

Update

THIS PAGE BLANK (USPTO)

Figure 7



THIS PAGE BLANK (USPTO)





Figure 8

THIS PAGE BLANK (USPTO)

THIS PAGE BLANK (USPTO)

Figure 10



THIS PAGE BLANK (USPTO)

PT/GB99/01772

BT GROUP LEGAL SERVICES

4/6/99

BEST AVAILABLE COPY

THIS PAGE BLANK (USPTO)